(De-)Composing Causality in Labeled Transition Systems

Georgiana Caltais
Department for Computer and Information Science
University of Konstanz, Germany
georgiana.caltais@uni-konstanz.de

Stefan Leue
Department for Computer and Information Science
University of Konstanz, Germany
stefan.leue@uni-konstanz.de

Mohammad Reza Mousavi
Centre for Research on Embedded Systems
Halmstad University, Sweden
m.r.mousavi@hh.se

In this paper we introduce a notion of counterfactual causality in the Halpern and Pearl sense that is compositional with respect to the interleaving of transition systems. The formal framework for reasoning on what caused the violation of a safety property is established in the context of labeled transition systems and Hennessy Milner logic. The compositionality results are devised for non-communicating systems.

1 Introduction

Determining and computing causalities is a frequently addressed issue in the philosophy of science and engineering, for instance when causally relating system faults to system failures. A notion of causality that is frequently used in relation to technical systems relies on counterfactual reasoning. Lewis [20] formulates the counterfactual argument, which defines when an event is considered a cause for some effect, in the following way: a) whenever the event presumed to be a cause occurs, the effect occurs as well, and b) when the presumed cause does not occur, the effect will not occur either (counterfactual argument). Counterfactual reasoning hence requires the consideration of alternative worlds: one world, corresponding to one program or system execution in software and systems analysis, where both the cause and the effect occur, and another world in which neither the cause nor the effect occur. Cause and effect are assumed to be temporally ordered.

In their seminal paper [13], Halpern and Pearl argue that the simple Lewis-style counterfactual argument cannot explain causalities if the causes correspond to complex logical structures of multiple events. Halpern and Pearl define a notion of complex logical events based on boolean equation systems and propose a number of conditions, called actual cause (AC) conditions, under which an event can be considered causal for an effect. The AC conditions encompass a couterfactual argument.

The Halpern and Pearl model of actual causation has been related in various forms to computing systems. Most relevant for our work is the work on causality checking [18, 17] which interprets the Halpern and Pearl event model and notion of actual causation in the context of the transition system and trace model for concurrent system computations. In addition to the Halpern and Pearl model, in causality checking the order of events as well as the non-occurrence of events can be causal. An implementation of causality checking using explicit-state model checking [19] as well as SAT-based bounded model checking [3] have been provided. The causality checking approach has been applied to various case studies in the area of analyzing critical systems for safety violations. In this setting, an ordered sequence of events is computed as being the actual cause of a safety property violation. In safety engineering the safety property violation is usually referred to as a hazard. The computed causalities will be displayed

G. Gössler and O. Sokolsky (Eds.): 1st Workshop on Causal Reasoning for Embedded and safety-critical Systems Technologies (CREST'16)
EPTCS 224, 2016, pp. 10–24, doi:10.4204/EPTCS.224.3

© G. Caltais, S. Leue & M.R. Mousavi
This work is licensed under the
Creative Commons Attribution License.

as fault trees complemented by temporal logic formulae which specify the order in which causal events occur.

The objective of this paper is to consider the notion of counterfactual causality reasoning and actual causation in the context of labeled transition systems (LTS's). In our setting the LTS's represent system models and Hennessy Milner logic (HML) [14] formulae specify the system properties for whose violation actual causes are sought. We also establish first results on computing causalities in this setting using (de-)compositional verification.

Our notion of causality complies to the characteristics of "actual causation" proposed in [13] and further adapted to the setting of concurrent systems in [17]. Intuitively, an execution within an LTS is causal whenever it leads to a state where a certain effect, or hazard, is enabled. We handle effects such as the violation of a safety property expressed in HML. Moreover, our definition includes a counterfactual test witnessing that a certain LTS execution L is causal for the occurrence of an effect E if and only if, were L not to happen, E would not occur either. Additionally, our definition exploits what is referred to as the "non-occurrence of events" in [17], and identifies relevant system execution fragments that, whenever performed, change the occurrence of the effect from true to false. Then, similarly to the approaches in [13, 17], our definition indicates that a setting that does not include the relevant executions discussed above has no influence on the effect as long as the causal events are present. Finally, we require causal executions to be minimal.

We establish the compositionality results with respect to the interleaving of LTS's, thus shifting the fault localization issue to the level of smaller interleaved components. The current approach only handles non-communicating LTS's. As an immediate extension of our approach, we would like to extend it to communicating LTS's by adopting ideas from [1, 8] (please see the conclusions section for more details on this extension).

Related work. Lewis-style counterfactual arguments have become the basis for a number of fault analysis, failure localization and software debugging techniques, such as delta debugging [26], nearest neighbor queries [23], counterexample explanation in model checking [12, 11] and why-because-analysis [15].

(De-)compositional verification has been studied in various contexts, such as model-checking [2, 6, 25] and model-based conformance testing [22, 24]. Our approach is based on our earlier work on decompositional verification of modal mu-calculus formulae [1]. Regarding compositional verification of causality, we are only aware of the line of work by Gößler, Le Métayer, and associates such as [9, 7, 8, 10]. In the remainder, we review [9] and [8] as two closely related examples in this line of work.

In [9], the authors define three trace-theoretic notions of causality for safety properties and provide an assume guarantee framework which allows for decomposing the identification of causes. They also provide decidability results. Their approach substantially differs from ours: firstly, we combine the different aspects of causality (positive causality, counterfactual, non-occurrence of events, and minimality) in one definition while in [9] a subset of these aspects is considered in three different definitions. Secondly, the approach of [9] relies on an assume-guarantee style of specifying the properties, with given LTS models for assume and guarantee contracts, while we rely on the alphabet of the system in decomposing the modal property and its cause. Our approach is in its early stages of development and the approach of [9] has been worked out in various directions. For example, [9] supports interaction models and is equipped with complexity and decidability results.

In [8], a de-compositional approach to a detecting a trace-based notion of causality is proposed. To start with a failed trace of the system, i.e., a counter-example of the property at hand, is consider and subsequently it is analyzed how the alternative possible behaviors of the different components may lead

to failed traces. In our approach, however, we do not start from a system-level counter-example: we aim at decomposing the modal formula for the property, so that all counter-examples are generated locally from the component specifications. Our initial results reported in this paper only concern interleaving components for which a very neat decomposition can be obtained, but our long-term vision is that modal decomposition will enable mechanized decomposition of the modal formula for communicating components, following the approach of [16, 1].

A trace-based approach to identifying causality for failures of interleaved systems has been recently introduced in [4]. In short, the authors propose a method for identifying event sequences that frequently occur within failing system executions, thus possibly revealing causes for system failures. One of the main differences with our approach is that in [4] system events are parameterised by thread identifiers, program and memory locations, while we consider more abstract events ranging over alphabets denoting (atomic) system actions. Nevertheless, the idea of using thread identifiers might be worth exploited in the context of extending our current work to the setting of concurrent, communicating LTS's.

Paper structure. In Section 2 we provide a brief reminder of HML, LTS's, and introduce LTS computations. In Section 3 we introduce our notion of causality and provide a series of examples motivating and explaining our definition. In Section 4 we discuss the (de-)compositionality results for causality. In Section 5 we conclude and provide pointers to further developments. For a more detailed version of this paper, including complete proofs of the compositionality results, we refer to [5].

2 Preliminaries

Let A be a possibly infinite set of labels, usually referred to as *alphabet*. Let $(-)^*$ be the Kleene star operator. We use $w, w_0, w_1, \ldots$ to range over words in A^*. We write ε for the empty word and wa for the word obtained by concatenating $w \in A^*$ and $a \in A$. We call a *sub-word* of a word w a word w' obtained by deleting n letters ($n \geq 1$) at some not-necessarily-adjacent positions in w, written $w' \in sub(w)$. The empty sequence ε is a sub-word of w.

Definition 1 (Labeled Transition Systems)**.** *A* labeled transition system *(LTS) is a triple* $(\mathbb{S}, s_0, A, \rightarrow)$*, where* $\mathbb{S}$ *is the set of states,* $s_0 \in \mathbb{S}$ *is the initial state,* A *is the action alphabet and* $\rightarrow \subseteq \mathbb{S} \times A \times \mathbb{S}$ *is the transition relation.*

We write $\twoheadrightarrow \subseteq \mathbb{S} \times A^* \times \mathbb{S}$, to denote the reachability relation, i.e., the smallest relation satisfying: $\frac{}{p \overset{\varepsilon}{\twoheadrightarrow} p}$, and $\frac{p \overset{w}{\twoheadrightarrow} p' \quad p' \overset{a}{\rightarrow} p''}{p \overset{wa}{\twoheadrightarrow} p''}$.

The set of actions that can be triggered as a first step from $s \in \mathbb{S}$ is denoted by $init(s)$: $init(s) = \{a \in A \mid \exists s' \in S \,:\, s \overset{a}{\rightarrow} s'\}$.

Definition 2 (Computations)**.** *Let* $[-]$ *be a list constructor. We write* $\mathscr{D} = [w_0, \ldots, w_n]$ *for a finite list of words* $w_i \in A^*$*, with* $0 \leq i \leq n$*. A notation of shape* $\mathscr{D} = [w_0, w_1, \ldots]$ *refers to an infinite list* $\mathscr{D}$ *of words* $w_i \in A^*$*, for* $i \geq 0$*. We write* $[\,]$ *to denote the empty list. Moreover, we write* $w : \mathscr{D}$ *as an alternative to a list with* w *as the first element, and* $\mathscr{D}$ *the "remaining" elements; for instance,* $w_1 : [w_2, w_3] = [w_1, w_2, w_3]$*. We say that lists* $\mathscr{D}_0, \ldots, \mathscr{D}_n$ *are* size-compatible *if they are finite lists of the same length, or if they are all infinite lists. For instance,* $[\,]$ *and* $[\,]$ *are size-compatible,* $[w_0, w_1, w_2]$ *and* $[w'_0, w'_1, w'_2]$ *are size-compatible,* $[w_0, w_1, \ldots]$ *and* $[w'_0, w'_1, \ldots]$ *are size-compatible, whereas* $[\,]$ *and* $[w]$ *are not size-compatible.*

Consider an LTS $T = (\mathbb{S}, s_0, A, \rightarrow)$ *and* $\pi \in (\mathbb{S} \times A \times [A^*])^* \times \mathbb{S}$ *a sequence*

$$(s_0, l_0, \mathscr{D}_0), \ldots (s_n, l_n, \mathscr{D}_n), s_{n+1}$$

over states $s_i \in \mathbb{S}$, actions $l_i \in A$ and sets of words $\mathscr{D}_i \subseteq A^*$, for $0 \leq i \leq n$. Whenever $\mathscr{D}_0, \ldots, \mathscr{D}_n$ are size-compatible, we write $traces((l_0, \mathscr{D}_0) \ldots (l_n, \mathscr{D}_n))$ or, in short, $traces(\pi)$, to denote the pairwise extensions of $l_0 \ldots l_n$ with words from $\mathscr{D}_0, \ldots, \mathscr{D}_n$ as follows:

$$\begin{array}{rcl} traces((l_0,[]) \ldots (l_n,[])) & = & \{l_0 \ldots l_n\} \\ traces((l_0, w_0 : \mathscr{D}_0) \ldots (l_n, w_n : \mathscr{D}_n)) & = & \{l_0 w_0 \ldots l_n w_n\} \cup traces((l_0, \mathscr{D}_0) \ldots (l_n, \mathscr{D}_n)) \end{array}$$

For instance, $traces((a,[w_{a0}, w_{a1}, w_{a2}]), (b,[\varepsilon, \varepsilon, \varepsilon]), (c,[\varepsilon, w_{c1}, \varepsilon])) = \{aw_{a0}bc, aw_{a1}bcw_{c1}, aw_{a2}bc\}$, for $a, b, c \in A$ and $w_{a0}, w_{a1}, w_{a2}, w_{c1} \in A^*$.

We say that π is a computation *of T whenever the following hold:*

- $s_0 \xrightarrow{l_0} s_1 \ldots \xrightarrow{l_n} s_{n+1}$,
- *$\mathscr{D}_0, \ldots, \mathscr{D}_n$ are size-compatible, and*
- *for all $w \in traces(\pi)$ there exists $s \in \mathbb{S}$ such that $s_0 \xRightarrow{w} s$.*

A computation consisting of only one state s_0 is called trivial computation. *We use $\pi, \mu, \ldots$ to range over computations.*

The set of sub-computations *of $\pi = (s_0, l_0, \mathscr{D}_0), \ldots, (s_n, l_0, \mathscr{D}_n), s_{n+1}$, denoted by $sub(\pi)$ is the set of all computations $\pi' = (s_0, l'_0, \mathscr{D}'_0), \ldots, (s_m, l'_m, \mathscr{D}'_m), s'_{m+1}$ such that $l'_0 \ldots l'_m \in sub(l_0 \ldots l_n)$. Note that all elements of $sub(\pi)$ should be computations themselves.*

For an intuition, size-compatible lists $\mathscr{D}_0, \ldots, \mathscr{D}_n$ encode the pairwise extensions of execution traces $l_0 \ldots l_n$ in T that always disable a certain effect. Given a computation $(s_0, l_0, \mathscr{D}_0), \ldots, (s_n, l_n, \mathscr{D}_n), s_{n+1}$ as above, sequences $w = l_0 w_0 \ldots l_n w_n \in traces((l_0, \mathscr{D}_0) \ldots (l_n, \mathscr{D}_n))$ determine executions $s_0 \xRightarrow{w} s$ in T, such that the effect does not occur in s. In our framework, occurrence of effects is formalised in terms of satisfiability of formulae in Hennessy Milner logic [14].

Definition 3 (Hennessy-Milner logic)**.** *The syntax of Hennessy-Milner logic (HML) [14] is given by the following grammar:*

$$\phi, \psi ::= \top \mid \langle a \rangle \phi \mid [a]\phi \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \qquad (a \in A).$$

We define the satisfaction relation $\models$ over LTS's and HML formulae as follows. The alphabet of a formula ϕ, denoted by $alphabet(\phi)$ is the set of actions that appear in ϕ.

Let $T = (\mathbb{S}, s_0, A, \rightarrow)$ be an LTS. Let ϕ, ϕ' range over HML formulae. It holds that:

$s \models \top$ for all $s \in \mathbb{S}$
$s \models \neg\phi$ whenever s does not satisfy ϕ; also written as $s \not\models \phi$
$s \models \phi \wedge \phi'$ if and only if $s \models \phi$ and $s \models \phi'$
$s \models \phi \vee \phi'$ if and only if $s \models \phi$ or $s \models \phi'$
$s \models \langle a \rangle \phi$ if and only if $s \xrightarrow{a} s'$ for some $s' \in \mathbb{S}'$ such that $s' \models \phi$
$s \models [a]\phi$ if and only if $s' \models \phi$ for all $s' \in \mathbb{S}'$ such that $s \xrightarrow{a} s'$.

3 Defining Causality

We further provide a notion of causality for LTS's. The effects that we consider are safety properties expressed as HML formulae. Examples motivating and explaining each of the items of our definition are given towards the end of this section.

Our notion of causality complies with that of "actual causation" proposed in [13] and further adapted to the setting of concurrent systems in [17]:

- Intuitively, AC1 in Definition 4 states that there must be a setting, or an execution within the LTS under consideration, that determines an effect, or a hazardous situation in which a safety property is violated.
- AC2(a) identifies a setting in which the effect does not occur. This is the counter-factual part of our definition.
- AC2(b) indicates that, as long as the causal events are present, a setting that does not include the relevant executions discussed above has no influence on the effect.
- AC2(c) corresponds to the so-called "non-occurrence of events" in [17], and identifies relevant system execution fragments that, whenever performed, change the occurrence of the effect from true to false. Intuitively, the aforementioned execution fragments are causal by their absence: the effect is enabled only within settings in which the fragments are not executed by our LTS.
- AC3 corresponds to the minimality condition in both [13] and [17].

The approach in [17] also exploits an ordering condition (OC) that identifies whether the order in which certain events are executed is causal with respect to a given effect, or not. Our framework does not explicitly handle such orderings. Nevertheless, for non-interleaved systems, such orderings are implicitly captured by sequences $l_0 \dots l_n$ determined by causal computation as in Definition 4. Additionally, as also discussed in Remark 1, the compositionality results in Section 4 can alleviate the ordering issue for certain kinds of effects in the context of interleaved systems.

Definition 4 (Causality for LTS's)**.** *Consider a transition system $T = (\mathbb{S}, s_0, A, \rightarrow)$; causal traces for an HML property ϕ in T denoted by $Causes(\phi, T)$ is the set of all computations $\pi = (s_0, l_0, \mathscr{D}_0), \dots, (s_n, l_n, \mathscr{D}_n), s_{n+1} \in (S \times A \times [A^*])^* \times S$ such that*

1. $s_0 \xrightarrow{l_0} \dots s_n \xrightarrow{l_n} s_{n+1} \wedge s_{n+1} \vDash \phi$ ***(Positive causality, AC1)****,*
2. $\exists \chi \in A^*, s' \in \mathbb{S} : s_0 \xRightarrow{\chi} s' \wedge s' \vDash \neg\phi$ ***(Counter-factual, AC2(a))****,*
3. $\forall \chi' = l_0\chi_0 \dots l_n\chi_n \in \{l_0 \dots l_n\} \cup (A^* \setminus traces((l_0, \mathscr{D}_0) \dots (l_n, \mathscr{D}_n))), s' \in \mathbb{S} : s_0 \xRightarrow{\chi'} s' \Rightarrow s' \vDash \phi$ ***(Causality of occurrence, AC2(b))***
4. $\forall \chi' \in traces((l_0, \mathscr{D}_0) \dots (l_n, \mathscr{D}_n)) \setminus \{l_0 \dots l_n\}, s' \in \mathbb{S} : s_0 \xRightarrow{\chi'} s' \Rightarrow s' \vDash \neg\phi$ ***(Causality of non-occurrence, AC2(c))***
5. $\forall \pi' \in sub(\pi) : \pi'$ *does not satisfy items 1. – 4. above* ***(Minimality, AC3)***

Definition 5 (Causal projection)**.** *A causal projection of $T = (\mathbb{S}, s_0, A, \rightarrow)$ with respect to an HML property ϕ, is $T' = (\mathbb{S}', s_0, A, \rightarrow')$ such that $\mathbb{S}' = \{s_i \mid 0 \leq i \leq n+1 \wedge (s_0, l_0, \mathscr{D}_0), \dots, (s_n, l_n, \mathscr{D}_n), s_{n+1} \in Causes(\phi, T)\}$ and $\rightarrow' = \{(s_i, l_i, s_{i+1}) \mid 0 \leq i \leq n \wedge (s_0, l_0, \mathscr{D}_0), \dots, (s_n, l_n, \mathscr{D}_n), s_{n+1} \in Causes(\phi, T)\}$.*

We write $T \downarrow \phi$ to denote the causal projection of T with respect to ϕ.

Intuitively, a causal projection is an LTS whose executions capture precisely all causal sequences determined by computations as in Definition 4.

Next, we illustrate the different aspects of Definition 4 using the following small "canonical" examples. The first example below motivates the positive causality condition (item 1 in Definition 4).

Example 1 (Positive causality)**.** *Consider the formula $\phi = \langle h \rangle \top$, which states that action h (for hazard) is enabled at the current state and LTS T_1 depicted in Figure 1.(a).*

The intuition behind the notion of cause suggests that action a should be considered a cause for $\langle h \rangle \top$. According to Definition 4, we have that $(s_{10}, a, [h]), s_{11} \in Causes(\phi, T)$. The causal projection of T_1 for ϕ is has one transition, namely, $s_{10} \xrightarrow{a} s_{11}$.

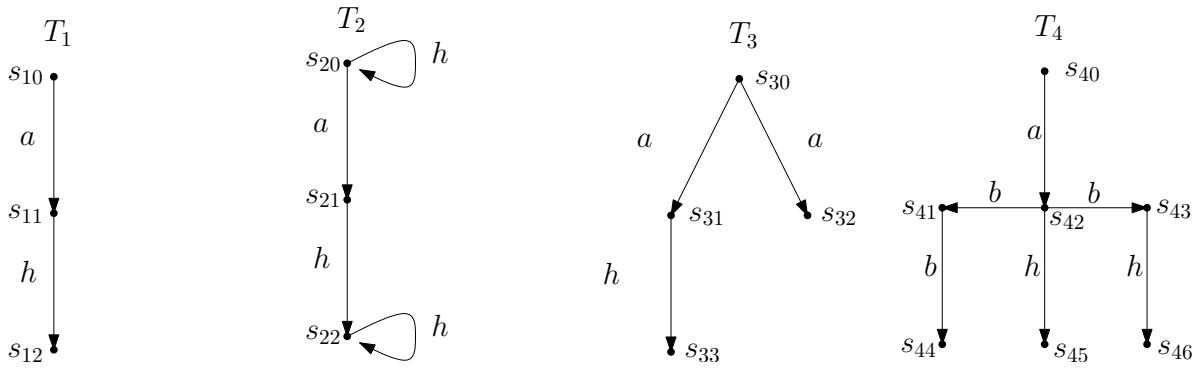

(a) Action *a* causes hazard *h*.

(b) The occurrence of hazard *h* is factual (trivial).

(c) The occurrence of *a* is not causal for hazard *h*.

(d) The non-occurrence of *bb* is causal for hazard *h*.

Figure 1: Canonical examples motivating different conditions on causality

The following example motivates the non-triviality condition (item 2 in Definition 4).

Example 2 (Counter-factual)**.** *Consider the LTS T_2 depicted in Figure 1.(b) and the same formula $\phi = \langle h \rangle \top$. Although trace a can lead to a state where ϕ holds, the hazard formula holds trivially everywhere else, and hence there is no cause to be identified; we refer to Lemma 1 for a formalisation.*

The next two examples motivate the causality of occurrence and non-occurence, respectively (items 3 and 4 in Definition 4).

Example 3 (Causality of Occurrence)**.** *Consider the LTS T_3 depicted in Figure 1.(c) and the same formula $\phi = \langle h \rangle \top$. Trace a can non-deterministically lead to two states, namely s_{31} and s_{32}. The formula holds only in one of them, namely in s_{31}. Hence, a cannot be considered a cause for the hazard. More precisely, if a trace is causal then its execution, or "occurrence", always leads to a state where the hazard holds.*

Example 4 (Causality of Occurrence and Non-occurrence)**.** *Consider the LTS T_4 depicted in Figure 1.(d) and the same formula $\phi = \langle h \rangle \top$. Trace a leads to state s_{42} where the hazard formula holds. Trace ab also leads to a hazardous state s_{43}; however, performing another b, i.e., performing the trace abb from the initial state, removes the hazard. Hence, $(s_{40}, a, [\varepsilon]), s_{42}$ is not in the set of causes for ϕ, because extending a with bb, for instance, violates ϕ and thereby violating item 3 in Definition 4. However, $(s_{40}, a, [h, bb, bh]), s_{42}$ is a cause, because a leads to a hazard, all possible extensions of a with anything but h, bb or bh, the only ones being ε and b, also keep the hazard. On the other hand, the extensions of a with h, bb or bh remove the hazard. Hence, h, bb and bh are the "relevant extension" that enable removing the hazard.*

The next example motivates the minimality condition, item 5 in Definition 4.

Example 5 (Minimality Condition)**.** *Consider again the LTS T_4 treated in Example 4. Computation $(s_{40}, a, [\varepsilon, \varepsilon]), (s_{42}, b, [h, b]), s_{43}$ is not a cause because it is not minimal (violating item 5 in Definition 4). This is because its sub-computation $(s_{40}, a, [h, bb, bh]), s_{42}$ is a cause as illustrated in Example 4.*

Consider the LTS T_5 depicted in Figure 2.(a) and the formula $\phi = \langle h \rangle \top$. For instance, the computation $(s_{50}, a, [\varepsilon, \varepsilon, \varepsilon \ldots]), (s_{51}, i, [h, ih, iih \ldots]), s_{51}$ is not in Causes(ϕ, T_5), because performing an i does not change the state of the system and hence, cannot contribute to the occurrence of the hazard. Computation $(s_{50}, a, [h, ih, iih, \ldots]), s_{51}$, however, is in Causes$(\phi, T_5)$, because it satisfies all the conditions of the cause, including minimality.

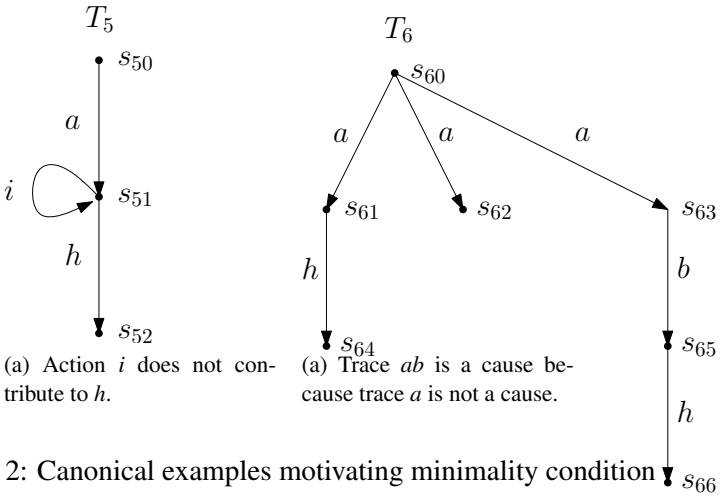

(a) Action i does not contribute to h.

(a) Trace ab is a cause because trace a is not a cause.

Figure 2: Canonical examples motivating minimality condition

Consider the LTS T_6 depicted in Figure 2.(b) and the formula $\phi = \langle h \rangle \top$. Computation $(s_{60}, a, [\varepsilon])$, $(s_{63}, b, [h])$, s_{65} is a cause for ϕ, despite the fact that computation $(s_{60}, a, [h, bh]), s_{61}$ also leads to the hazard.

This is not a violation of minimality, because $(s_{60}, a, [h, bh]), s_{61}$ does not satisfy the so-called "Causality of occurrence" (AC2(b)) in Definition 4, as also illustrated in Example 3.

4 (De-)composing Causality

In this section we provide the main results regarding (de-)compositionality of causality. Theorem 1 states the equivalence between reasoning on causality with respect to disjunctions $\phi \vee \psi$ of HML formulae in the context of interleaved LTS's, and reasoning on causality with respect to ϕ or ψ in the corresponding interleaved components. Orthogonally, Theorem 2 captures the equivalence between reasoning on causality with respect to conjunctions $\phi \wedge \psi$ of HML formulae in the context of interleaved LTS's, and reasoning on causality with respect to ϕ and ψ in the corresponding interleaved components. Both results are established for non-communicating LTS's executing disjoint sets of actions.

Our formal framework exploits standard notions of interleaving ($||$) and non-deterministic ($+$) choice between LTS's [21] or, more explicitly, between causal projections as in Definition 5. Consider the LTS $T = (\mathbb{S}, s_0, A, \rightarrow)$, $a \in A$ and $s, s', p, p' \in \mathbb{S}$. Then:

$$s \,||\, p \xrightarrow{a} s' \,||\, p \;\text{ whenever }\; s \xrightarrow{a} s' \qquad s + p \xrightarrow{a} s' \;\text{ whenever }\; s \xrightarrow{a} s'$$
$$s \,||\, p \xrightarrow{a} s \,||\, p' \;\text{ whenever }\; p \xrightarrow{a} p' \qquad s + p \xrightarrow{a} p' \;\text{ whenever }\; p \xrightarrow{a} p'.$$

Consider LTS's $T = (\mathbb{S}, s_0, A, \rightarrow)$ and $T' = (\mathbb{S}', s'_0, B, \rightarrow')$. We abuse the notation and write $T||T'$ in lieu of $s_0 \,||\, s'_0$, and $T + T'$ in lieu of $s_0 + s'_0$.

With this intuition in mind, we proceed to discussing our compositionality results.

Lemma 1 provides a result that shows that reasoning on (de-)composition of causality in the context of formulae that hold in the initial state of a system is trivial.

Lemma 1 (Immediate Causality)**.** *Consider the LTS's $T = (\mathbb{S}, s_0, A, \rightarrow)$ and the HML property ϕ. If $s_0 \vDash \phi$ it holds that $s_0 = Causes(\phi, T)$ or $Causes(\phi) = \emptyset$.*

We call properties ϕ as above *immediate effects*.

4.1 (De-)Composing Disjunction

In what follows we show that reasoning on causality with respect to disjunctions of HML formulae $\phi \vee \psi$ can be performed in a compositional fashion.

Intuitively, the result in Lemma 2 states that causality is preserved under disjunction of HML formulae and the interleaving of non-communicating LTS's. Or, more precisely, given two non-communicating LTS's T and T' and two HML formulae ϕ and ψ built over their corresponding alphabets, it holds that a cause $\pi \in \mathit{Causes}(\phi, T)$ determines a cause $\mu \in \mathit{Causes}(\phi \vee \psi, T \mid\mid T')$ within the interleaved LTS's.

Lemma 2. *Consider LTS's $T = (\mathbb{S}, s_0, A, \rightarrow)$ and $T' = (\mathbb{S}', s'_0, B, \rightarrow')$ such that $A \cap B = \emptyset$. Assume two HML formulae ϕ and ψ over A and B, respectively. Whenever ϕ and ψ are not immediate effects, the following holds:*

If $\pi = (s_0, l_0, \mathscr{D}_0), \ldots, (s_n, l_n, \mathscr{D}_n), s_{n+1} \in \mathit{Causes}(\phi, T)$, then there exists
$\mu = (s_0 \mid\mid s'_0, l_0, \overline{\mathscr{D}}_0), \ldots, (s_n \mid\mid s'_0, l_n, \overline{\mathscr{D}}_n), s_{n+1} \mid\mid s'_0 \in \mathit{Causes}(\phi \vee \psi, T \mid\mid T')$.

Proof Sketch. The statement follows by two intermediate results.

We show how to create a computation μ satisfying conditions AC1–AC2(c) in Definition 4 from π, given the hypothesis that π satisfies conditions AC1–AC2(c) as well. AC1 is satisfied for μ as a consequence of AC1 being satisfied for π. AC2(a) trivially holds for μ as ϕ and ψ are not immediate effects. Showing AC2(b) and AC2(c) strongly relies on the shape of $\overline{\mathscr{D}}_0, \ldots, \overline{\mathscr{D}}_n$. The lists $\overline{\mathscr{D}}_i$ are created in three steps.

1. We begin by simply "copying" the information in each $\mathscr{D}_i$ into the corresponding $\overline{\mathscr{D}}_i$.
2. We identify all causal traces χ obtained by interleaving the causal traces of π with the causal traces determined by all computations in $\mathit{Causes}(\psi, T')$. We make the necessary insertions into the lists $\overline{\mathscr{D}}_i$, so that χ's are stored as causal traces of computations in $\mathit{Causes}(\phi \vee \psi, T \mid\mid T')$.
3. We compute all the causal traces χ for $\phi \vee \psi$ that do not allow s'_0 to evolve in T', but consist of words in B as well. We make the necessary insertions into the lists $\overline{\mathscr{D}}_i$, so that χ's are stored as causal traces of computations in $\mathit{Causes}(\phi \vee \psi, T \mid\mid T')$. This step guarantees that the remaining traces in $(A \cup B)^* \setminus \mathit{traces}((l_0, \overline{\mathscr{D}}_0) \ldots ((l_n, \overline{\mathscr{D}}_n)))$ are not "harmful" with respect to AC2(b) for μ, as they never lead to $s \mid\mid s' \vDash \neg\phi \wedge \neg\psi$.

By the above construction, AC2(b) and AC2(c) hold for μ as well.

AC3 for μ is proved to hold by reductio ad absurdum. In short, we show that whenever there is $\mu' \in \mathit{sub}(\mu)$, such that μ' satisfies AC1–AC2(b), there exists $\pi' \in \mathit{sub}(\pi)$, such that π' satisfies AC1–AC2(b) as well. This contradicts the hypothesis $\pi \in \mathit{Causes}(\phi, T)$.

□

Intuitively, Lemma 3 states that causality with respect to an effect $\phi \vee \psi$ in two interleaved, but non-communicating LTS's, is preserved by at least one of the interleaved components. Or, more precisely, given two non-communicating LTS's T and T' and two HML formulae ϕ and ψ built over their corresponding alphabets, it holds that a cause $\mu \in \mathit{Causes}(\phi \vee \psi, T \mid\mid T')$ within the interleaved LTS's determines a cause $\pi \in \mathit{Causes}(\phi, T)$ for ϕ in T, or a cause $\pi' \in \mathit{Causes}(\psi, T')$ for ψ in T'.

Lemma 3. *Consider LTS's $T = (\mathbb{S}, s_0, A, \rightarrow)$ and $T' = (\mathbb{S}', s'_0, B, \rightarrow')$ such that $A \cap B = \emptyset$. Assume two HML formulae ϕ and ψ over A and B, respectively. Whenever ϕ and ψ are not immediate effects, the following holds:*

If $\mu = (s_0 \mid\mid s'_0, l_0, \mathscr{D}_0), \ldots, (s_n \mid\mid s'_n, l_n, \mathscr{D}_n), s_{n+1} \mid\mid s'_{n+1} \in \mathit{Causes}(\phi \vee \psi, T \mid\mid T')$*, then there exists*
$\pi = (s_k, l_k, \overline{\mathscr{D}_k}), \ldots, (s_m, l_m, \overline{\mathscr{D}_m}), s_{n+1} \in \mathit{Causes}(\phi, T)$ *or*
$\pi' = (s'_p, l'_p, \overline{\mathscr{D}'_p}), \ldots, (s'_q, l'_q, \overline{\mathscr{D}'_q}), s'_{n+1} \in \mathit{Causes}(\psi, T')$.

For all $k \leq i \leq m$*:* $(s_i, l_i, \overline{\mathscr{D}_i})$ *corresponds to* $(s_i \mid\mid s'_i, l_i, \mathscr{D}_i)$ *in* μ*, whenever* $l_i \in A$*. For all* $p \leq j \leq q$*:* $(s'_j, l'_j, \overline{\mathscr{D}'}_j)$ *corresponds to* $(s_j \mid\mid s'_j, l'_j, \mathscr{D}'_j)$ *in* μ*, whenever* $l'_j \in B$*. Moreover,* $l_k \ldots l_m = l_0 \ldots l_n \downarrow A$*,* $l'_p \ldots l'_q = l_0 \ldots l_n \downarrow B$.

Proof Sketch. The statement follows by two intermediate results.

First, we show that one can build π or π' as above, such that π or π' satisfy conditions AC1–AC2(c) in Definition 4, given the hypothesis that μ satisfies AC1–AC2(c) as well. The reasoning for proving this intermediate result strongly relies on the shape of the lists $\overline{\mathscr{D}}_i$ and $\overline{\mathscr{D}}'_j$ corresponding to π and π', respectively. We construct the aforementioned lists in three steps.

1. We start with empty lists $\overline{\mathscr{D}}_i$ and $\overline{\mathscr{D}}'_j$.
2. Then, we "encode" causal sequences $\chi \in \mathit{traces}((l_0, \mathscr{D}_0) \ldots (l_n, \mathscr{D}_n)) \setminus \{l_0 \ldots l_n\}$ satisfying AC2(c) by definition, into $\mathit{traces}((l_k, \overline{\mathscr{D}}_k) \ldots (l_m, \overline{\mathscr{D}}_m))$ and, respectively, $\mathit{traces}((l'_p, \overline{\mathscr{D}}'_p) \ldots (l'_q, \overline{\mathscr{D}}'_q))$, via the projections of χ on A and, respectively, B that satisfy AC2(c) as well.
3. Eventually, we "prepare" π for satisfying AC2(b). We identify all sequences $\chi \in A^* \setminus \mathit{traces}((l_k, \overline{\mathscr{D}}_k) \ldots (l_m, \overline{\mathscr{D}}_m))$ that always lead to $s \vDash \neg\phi$. For each such χ we make the necessary insertions into the lists $\overline{\mathscr{D}}_i$, so that χ's are stored as causal traces of computations in $\mathit{Causes}(\phi, T)$. We repeat the "preparation" process for π' as well.

Then, we show that π or π' satisfy AC1–AC2(c) by reductio ad absurdum. Without loss of generality, assume that π satisfies AC1–AC2(c). Showing that π has to satisfy *AC*3 as well follows by proof by contradiction. More explicitly, we show that whenever there exists $\widetilde{\pi} \in \mathit{sub}(\pi)$ satisfying AC1–AC2(c), one can construct $\widetilde{\mu} \in \mathit{sub}(\mu)$ such that $\widetilde{\mu}$ satisfies AC1–AC2(c) as well. This contradicts the hypothesis $\mu \in \mathit{Causes}(\phi \vee \psi, T \mid\mid T')$.

□

Corollary 1 states that a causal computation μ with respect to an effect $\phi \vee \psi$ in interleaved, but non-communicating LTS's, determines a causal computation π in the interleaved component that triggered the first step in μ.

Corollary 1. *Consider LTS's* $T = (\mathbb{S}, s_0, A, \rightarrow)$ *and* $T' = (\mathbb{S}', s'_0, B, \rightarrow')$ *such that* $A \cap B = \emptyset$*. Assume two HML formulae* ϕ *and* ψ *over A and B, respectively. Whenever* ϕ *and* ψ *are not immediate effects, the following holds:*

If $\mu = (s_0 \mid\mid s'_0, l_0, \mathscr{D}_0), \ldots, (s_n \mid\mid s'_n, l_n, \mathscr{D}_n), s_{n+1} \mid\mid s'_{n+1} \in \mathit{Causes}(\phi \vee \psi, T \mid\mid T')$ *then*

- *if* $l_0 \in A$ *then exists* $\pi = (s_k, l_k, \overline{\mathscr{D}_k}), \ldots, (s_m, l_m, \overline{\mathscr{D}_m}), s_{n+1} \in \mathit{Causes}(\phi, T)$*; otherwise*
- *if* $l_0 \in B$ *then exists* $\pi' = (s'_p, l'_p, \overline{\mathscr{D}'_p}), \ldots, (s'_q, l'_q, \overline{\mathscr{D}'_q}), s'_{n+1} \in \mathit{Causes}(\psi, T')$.

For all $k \leq i \leq m$*:* $(s_i, l_i, \overline{\mathscr{D}_i})$ *corresponds to* $(s_i \mid\mid s'_i, l_i, \mathscr{D}_i)$ *in* μ*, whenever* $l_i \in A$*. For all* $p \leq j \leq q$*:* $(s'_j, l'_j, \overline{\mathscr{D}'}_j)$ *corresponds to* $(s_j \mid\mid s'_j, l'_j, \mathscr{D}'_j)$ *in* μ*, whenever* $l'_j \in B$*. Moreover,* $l_k \ldots l_m = l_0 \ldots l_n \downarrow A$*,* $l'_p \ldots l'_q = l_0 \ldots l_n \downarrow B$.

Proof. The result follows immediately by Lemma 3, Lemma 2 and the minimality condition AC3 in Definition 4. □

Lemma 4 states that, as a consequence of the minimality condition, causal computations with respect to effects $\phi \vee \psi$ in interleaved, non-communicating LTS's capture executions of only one of the interleaved components.

Lemma 4. *Consider LTS's $T = (\mathbb{S}, s_0, A, \rightarrow)$ and $T' = (\mathbb{S}', s'_0, B, \rightarrow')$ such that $A \cap B = \emptyset$. Assume two HML formulae ϕ and ψ over A and B, respectively. Whenever ϕ and ψ are not immediate effects and $\mu \in Causes(\phi \vee \psi, T \,||\, T')$, then either*

$$\mu = (s_k \,||\, s'_0, l_k, \mathscr{D}_k), \ldots, (s_m \,||\, s'_0, l_m, \mathscr{D}_m), s_{n+1} \,||\, s'_0, \text{ or}$$
$$\mu = (s_0 \,||\, s'_p, l'_p, \mathscr{D}'_p), \ldots, (s_0 \,||\, s'_q, l'_q, \mathscr{D}'_q), s_0 \,||\, s'_{n+1}$$

such that, for all $k \leq i \leq m$ and $p \leq j \leq q$: $s_i \in \mathbb{S}$, $s'_j \in \mathbb{S}'$, $l_i \in A$, $l'_j \in B$, $\mathscr{D}_i \in A^*$ and $\mathscr{D}'_j \in B^*$.

Proof. Assume $\mu = (s_0 \,||\, s'_0, l_0, \overline{\mathscr{D}}_0), \ldots, (s_n \,||\, s'_n, l_n, \overline{\mathscr{D}}_n), s_{n+1} \,||\, s'_{n+1} \in Causes(\phi \vee \psi, T \,||\, T')$. Assume, without loss of generality, that by Lemma 3 there exists a computation:

$$\widetilde{\pi} = (s_k, l_k, \widetilde{\mathscr{D}}_k), \ldots, (s_m, l_m, \widetilde{\mathscr{D}}_m), s_{n+1} \in Causes(\phi, T)$$

such that for all $k \leq i \leq m$: $(s_i, l_i, \widetilde{\mathscr{D}}_i)$ corresponds to $(s_i \,||\, s'_i, l_i, \overline{\mathscr{D}}_i)$ in μ, whenever $l_i \in A$. Moreover, $l_k \ldots l_m = l_0 \ldots l_n \downarrow A$. Then, by Lemma 2, it follows that there exists a computation

$$\widehat{\mu} = (s_k \,||\, s'_0, l_k, \widehat{\mathscr{D}}_k), \ldots, (s_m \,||\, s'_0, l_m, \widehat{\mathscr{D}}_m), s_{n+1} \,||\, s'_0 \in Causes(\phi \vee \psi, T \,||\, T').$$

Additionally, observe that $\widehat{\mu} \in sub(\mu)$. This violates the minimality condition AC3 for μ, unless $\mu = \widehat{\mu}$. This proves our initial statement. □

Theorem 1 is the main result of this section. Intuitively, it states that reasoning on causality with respect to an effect $\phi \vee \psi$ in the context of non-communicating, interleaved LTS's is equivalent to reasoning on causality for ϕ or ψ in the context of the corresponding interleaved components.

Theorem 1 ((De-)composing Disjunction)**.** *Consider LTS's $T = (\mathbb{S}, s_0, A, \rightarrow)$ and $T' = (\mathbb{S}', s'_0, B, \rightarrow')$ such that $A \cap B = \emptyset$. Assume two HML formulae ϕ and ψ over A and B, respectively. Whenever ϕ and ψ are not immediate effects, the following holds:*

$$T \,||\, T' \downarrow (\phi \vee \psi) \;\simeq\; T \downarrow \phi + T' \downarrow \psi. \tag{1}$$

Proof. Let $(\mathbb{S}_{||}, s_0 \,||\, s'_0, A \cup B, \rightarrow_{||}) = (T \,||\, T') \downarrow (\phi \vee \psi)$ and $(\mathbb{S}_+, s_0 + s'_0, A \cup B, \rightarrow_+) = (T \downarrow \phi) + (T' \downarrow \psi)$, respectively. The result follows immediately by Corollary 1, Lemma 4 and the semantics of the non-deterministic choice operator $(+)$, where the isomorphic structure is underlined by:

$$\begin{array}{rcl} & f : \mathbb{S}_{||} \rightarrow \mathbb{S}_+ & \\ f(s_0 \,||\, s'_0) & = & \underline{s_0 + s'_0} \\ f(p \,||\, q) & = & \left\{ \begin{array}{ll} p & \text{if} \quad q = s'_0 \wedge p \neq s_0 \\ q & \text{if} \quad p = s_0 \wedge q \neq s'_0 \end{array} \right. \end{array} \qquad \begin{array}{rcl} & f^{-1} : \mathbb{S}_+ \rightarrow \mathbb{S}_{||} & \\ f^{-1}(s_0 + s'_0) & = & \underline{s_0 \,||\, s'_0} \\ f^{-1}(p) & = & \left\{ \begin{array}{ll} p \,||\, s'_0 & \text{if} \quad p \in \mathbb{S} \wedge p \neq s_0 \\ s_0 \,||\, p & \text{if} \quad p \in \mathbb{S}' \wedge p \neq s'_0 \end{array} \right. \end{array}$$

□

Example 6. *For an example, consider two LTS's T and T' with initial states s_0 and p_0, respectively, depicted as in Figure 3. Let $\phi = \langle h \rangle \top$ and $\psi = \langle h' \rangle \top$ be two HML formulae. It is straightforward to see that $T \downarrow \phi$ is defined by the dotted transition $s_0 \overset{a}{\dashrightarrow} s_1$ in T, whereas $T' \downarrow \psi$ is $p_0 \overset{d}{\dashrightarrow} p_1 \overset{e}{\dashrightarrow} p_2$. The interleaving of T and T' is the LTS originating in $s_0 \,||\, p_0$ in Figure 3. At a closer look, one can see that $T \,||\, T' \downarrow (\phi \vee \psi)$ is the transition system defined by the dotted transitions $s_0 \,||\, p_0 \overset{a}{\dashrightarrow} s_1 \,||\, p_0$ and $s_0 \,||\, p_0 \overset{d}{\dashrightarrow} s_0 \,||\, p_1 \overset{e}{\dashrightarrow} s_0 \,||\, p_2$, which is obviously isomorphic with $T \downarrow \phi + T' \downarrow \psi$.*

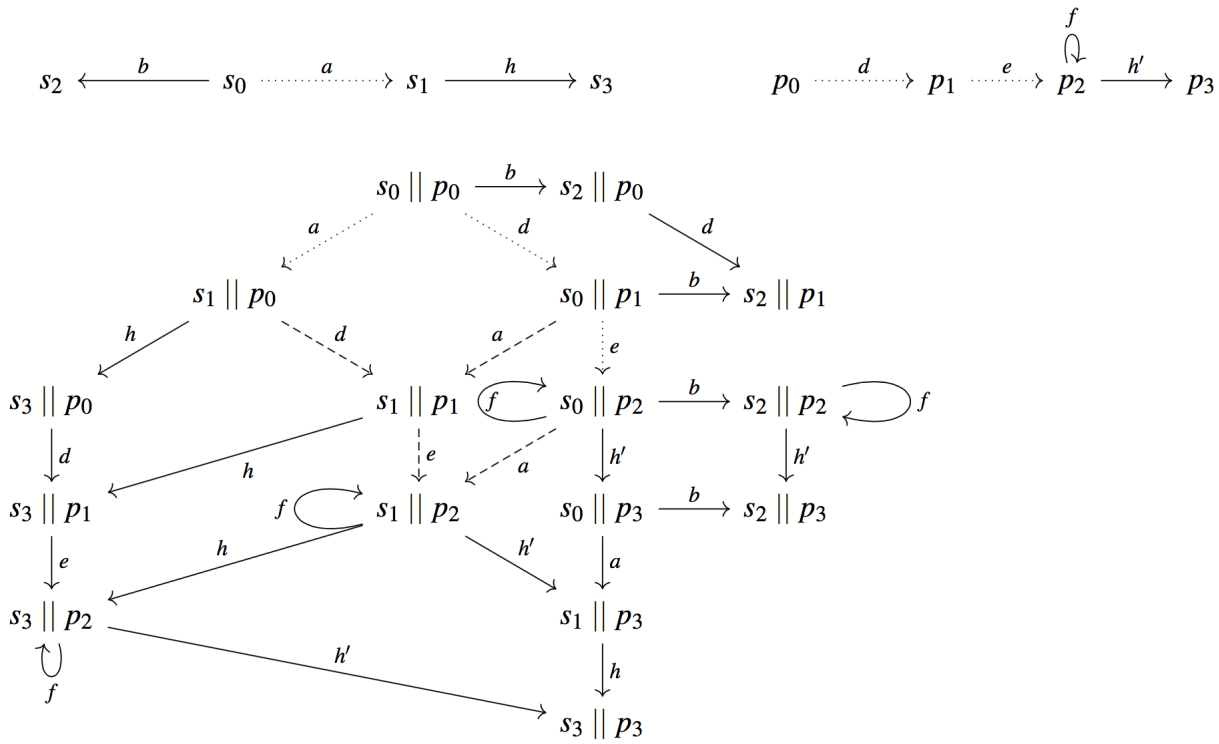

Figure 3: (De-)composing causality

4.2 (De-)Composing Conjunction

In what follows we show that reasoning on causality with respect to conjunctions of HML formulae $\phi \wedge \psi$ can be performed in a compositional fashion.

Lemma 5 states that causalities in two non-communicating LTS's are reflected within their interleaving as well.

Lemma 5. *Consider LTS's $T = (\mathbb{S}, s_0, A, \rightarrow)$ and $T' = (\mathbb{S}', s'_0, B, \rightarrow')$ such that $A \cap B = \emptyset$. Assume two HML formulae ϕ and ψ over A and B, respectively. Whenever ϕ and ψ are not immediate effects, the following holds. If*

$$\pi = (s_k, l_k, \mathscr{D}_k), \ldots, (s_m, l_m, \mathscr{D}_m), s_{m+1} \in \mathit{Causes}(\phi, T) \text{ and}$$
$$\pi' = (s'_p, l'_p, \mathscr{D}'_p), \ldots, (s'_q, l'_q, \mathscr{D}'_q), s'_{q+1} \in \mathit{Causes}(\psi, T') \text{ then}$$
$$\mu = (s_0 \,||\, s'_0, l_0, \overline{\mathscr{D}}_0), \ldots, (s_n \,||\, s'_n, l_n, \overline{\mathscr{D}}_n), s_{n+1} \,||\, s'_{n+1} \in \mathit{Causes}(\phi \wedge \psi, T \,||\, T')$$

for all μ such that $s_0 \,||\, s'_0 \xrightarrow{l_0} \ldots s_n \,||\, s'_n \xrightarrow{l_n} s_{n+1} \,||\, s'_{n+1}$ is an execution sequence in $s_k \xrightarrow{l_k} \ldots s_m \xrightarrow{l_m} s_{m+1} \,||\, s'_p \xrightarrow{l'_p} \ldots s'_q \xrightarrow{l'_q} s'_{q+1}$, and $s_0 \,||\, s'_0 = s_k \,||\, s'_p$, $s_n \,||\, s'_n = s_m \,||\, s'_q$, $s_{n+1} \,||\, s'_{n+1} = s_{m+1} \,||\, s'_{q+1}$, $l_0 \ldots l_n \downarrow A = l_k \ldots l_m$ and $l_0 \ldots l_n \downarrow B = l'_p \ldots l'_q$.

Proof Sketch. The statement is a consequence of two intermediate results.

First we show that whenever π and π' satisfy conditions AC1–AC2(c) in Definition 4, one can build μ as above, such that μ satisfies AC1–AC2(c) as well. Showing that μ satisfies AC1 and AC2 is immediate, by the assumption that both π and π' satisfy AC1–AC2(c) and the fact that ϕ and ψ are not immediate effects. Proving that AC2(b) and AC2(c) hold for μ strongly relies on the lists $\overline{\mathscr{D}}_i$ in μ. The construction of $\overline{\mathscr{D}}_i$'s is as follows.

1. We start with $\overline{\mathscr{D}}_i$'s set to the empty list [].
2. Then, note that all causal traces χ corresponding to π are causal for $\neg\phi \vee \neg\psi$ as well. Hence, we consider sequences $\overline{\chi}$ from the interleaving of such χ with $\chi' \in B^*$ and make the corresponding additions to all $\overline{\mathscr{D}}_i$'s, such that $\overline{\chi}$ is captured within $traces((l_0,\overline{\mathscr{D}}_0)\dots(l_n,\overline{\mathscr{D}}_n))$ as well. Symmetrically, repeat the procedure for all causal traces corresponding to π'.

   Intuitively, this step works also as a "cleaning" step preparing μ to satisfy AC2(b) w.r.t. $\phi \wedge \psi$.

At this point AC2(b) and AC2(c) hold for μ, by the construction of lists $\overline{\mathscr{D}}_i$ above.

Proving minimality of μ follows by reductio ad absurdum. The intuition is as follows. Whenever there exists $\mu' \in sub(\mu)$ such that μ' satisfies AC1–AC2(c), one can build $\widetilde{\pi} \in sub(\pi)$ and $\widetilde{\pi'} \in sub(\pi')$ such that $\widetilde{\pi}$ and $\widetilde{\pi'}$ satisfy AC1–AC2(c). This contradicts the hypothesis $\pi \in Causes(\phi, T)$ and $\pi' \in Causes(\psi, T')$.

□

Lemma 6 states that causality with respect to an HML formula $\phi \wedge \psi$ in the context of interleaved, non-communicating LTS's, determines causality with respect to ϕ and ψ in the corresponding interleaved components.

Lemma 6. *Consider LTS's $T = (\mathbb{S}, s_0, A, \rightarrow)$ and $T' = (\mathbb{S}', s'_0, B, \rightarrow')$ such that $A \cap B = \emptyset$. Assume two HML formulae ϕ and ψ over A and B, respectively. Whenever ϕ and ψ are not immediate effects, the following holds.*

If $\mu = (s_0 \,||\, s'_0, l_0, \overline{\mathscr{D}}_0), \dots, (s_n \,||\, s'_n, l_n, \overline{\mathscr{D}}_n), s_{n+1} \,||\, s'_{n+1} \in Causes(\phi \wedge \psi, T \,||\, T')$, then there exist

$$\pi = (s_k, l_k, \mathscr{D}_k), \dots, (s_m, l_m, \mathscr{D}_m), s_{m+1} \in Causes(\phi, T) \text{ and}$$
$$\pi' = (s'_p, l'_p, \mathscr{D}'_p), \dots, (s'_q, l'_q, \mathscr{D}'_q), s'_{q+1} \in Causes(\psi, T')$$

where $s_k \xrightarrow{l_k} \dots s_m \xrightarrow{l_m} s_{m+1} \,||\, s'_p \xrightarrow{l'_p} \dots s'_q \xrightarrow{l'_q} s'_{q+1}$ includes the execution sequence $s_0 \,||\, s'_0 \xrightarrow{l_0} \dots s_n \,||\, s'_n \xrightarrow{l_n} s_{n+1} \,||\, s'_{n+1}$, and $s_k \,||\, s'_p = s_0 \,||\, s'_0$, $s_m \,||\, s'_q = s_n \,||\, s'_n$, $s_{m+1} \,||\, s'_{q+1} = s_{n+1} \,||\, s'_{n+1}$, $l_k \dots l_m = l_0 \dots l_n \downarrow A$ and $l'_p \dots l'_q = l_0 \dots l_n \downarrow B$.

Proof Sketch. First, we show that one can build π and π' as above, such that π and π' satisfy conditions AC1–AC2(c) in Definition 4, given the hypothesis that μ satisfies AC1–AC2(c) as well. The reasoning for proving this intermediate result strongly relies on the shape of the lists $\mathscr{D}_i$ and $\mathscr{D}'_j$ corresponding to π and π', respectively. We construct the aforementioned lists in three steps.

1. We start with empty lists $\mathscr{D}_i$ and $\mathscr{D}'_j$.
2. Then, we "encode" causal sequences $\chi \in traces((l_0,\overline{\mathscr{D}}_0)\dots(l_n,\overline{\mathscr{D}}_n)) \setminus \{l_0 \dots l_n\}$ satisfying AC2(c) by definition, into $traces((l_k,\mathscr{D}_k)\dots(l_m,\mathscr{D}_m))$ and, respectively, $traces((l'_p,\mathscr{D}'_p)\dots(l'_q,\mathscr{D}'_q))$ as follows. Whenever χ always leads to states satisfying $\neg\phi$, make the corresponding additions to $\mathscr{D}_i$ such that the projection of χ on A is stored within $traces((l_k,\mathscr{D}_k)\dots(l_m,\mathscr{D}_m))$. Symmetrically, repeat the procedure for causal sequences χ that always lead to states satisfying $\neg\psi$.
3. Eventually, we "prepare" π for satisfying AC2(b). We identify all sequences $\chi \in A^* \setminus traces((l_k,\overline{\mathscr{D}}_k)\dots(l_m,\overline{\mathscr{D}}_m))$ that always lead to $s \vDash \neg\phi$. For each such χ we make the necessary insertions into the lists $\overline{\mathscr{D}}_i$, so that χ is stored as a causal trace of π. We repeat the "preparation" process for π' as well.

Then, we show that π and π' satisfy AC1–AC2(c) by reductio ad absurdum. Showing that π has to satisfy *AC*3 follows by proof by contradiction as well. Intuitively, we show that whenever there exists

$\widetilde{\pi} \in sub(\pi)$ satisfying AC1–AC2(c), one can construct $\widetilde{\mu} \in sub(\mu)$ such that $\widetilde{\mu}$ satisfies AC1–AC2(c) as well. This contradicts the hypothesis $\mu \in Causes(\phi \wedge \psi, T \mid\mid T')$. Similar reasoning for proving that π' has to satisfy $AC3$.

□

Theorem 2 is the main result of this section. Intuitively, it states that reasoning on causality with respect to an effect $\phi \wedge \psi$ in the context of non-communicating, interleaved LTS's is equivalent to reasoning on causality for ϕ and ψ in the context of the corresponding interleaved components.

Theorem 2 ((De-)composing Conjunction)**.** *Consider $T = (\mathbb{S}, s_0, A, \rightarrow)$ and $T' = (\mathbb{S}', s_0', B, \rightarrow')$ such that $A \cap B = \emptyset$. Assume two HML formulae ϕ and ψ over A and B, respectively. Whenever ϕ and ψ are not immediate effects, the following holds:*

$$T \mid\mid T' \downarrow (\phi \wedge \psi) \;=\; (T \downarrow \phi) \mid\mid (T' \downarrow \psi). \tag{2}$$

Proof. The result is immediate by Lemma 5 and Lemma 6. □

For an example, we refer again to the LTS's in Figure 3. The causal projection $T \mid\mid T' \downarrow (\phi \wedge \psi)$ is defined by the dashed/dotted transitions $s_0 \mid\mid p_0 \overset{d}{\dashrightarrow} s_0 \mid\mid p_1 \text{ - } \underline{a} \rightarrow s_1 \mid\mid p_1 \text{ - } \underline{e} \rightarrow s_1 \mid\mid p_2$, $s_0 \mid\mid p_0 \overset{d}{\dashrightarrow} s_0 \mid\mid p_1 \overset{e}{\dashrightarrow} s_0 \mid\mid p_2 \text{ - } \underline{a} \rightarrow s_1 \mid\mid p_2$ and $s_0 \mid\mid p_0 \overset{a}{\dashrightarrow} s_1 \mid\mid p_0 \text{ - } \underline{d} \rightarrow s_1 \mid\mid p_1 \text{ - } \underline{e} \rightarrow s_1 \mid\mid p_2$. This is precisely the interleaving of the causal projections $T \downarrow \phi$ and $T' \downarrow \psi$.

Remark 1. *As pointed out in Section 3, the proposed notion of causality does not check whether the order in which certain actions are executed is causal with respect to the violation of a safety property, or not. Nevertheless, as already mentioned, for non-interleaved systems such orderings are implicitly captured by sequences $l_0 \ldots l_n$ determined by causal computations as in Definition 4. Additionally, in the context of interleaved systems, the ordering information can be irrelevant. For formulae defined over disjoint alphabets, based on the compositionality results in Theorem 1 and Theorem 2, causal reasoning is "pushed" at the level of the interleaved components, hence the order in which these components execute the interleaving does not matter.*

5 Conclusions and Future Work

In this paper we introduce a notion of causality for LTS's and violation of safety properties expressed in terms of HML formulae. The proposed notion of causality inherits the characteristics of "actual causation" proposed in [13, 17] and, in addition, is compositional with respect to the interleaving of the considered type of non-communicating LTS's.

A natural extension is handling causality in the context of communicating LTS's in the style of CCS [21], for instance. The challenge would be to establish (de-)compositionality results whenever the interleaved systems display internal, non-observable behaviour. The current approach relies on the fact that the HML formulae are defined over "observable", disjoint alphabets. However, the general modal decomposition theorems such as those proposed in [16, 1] do provide support for arbitrary formulae and silent actions. This provides an interesting ground to extend our approach to communicating processes.

Of equal importance is extending our framework to handle causality for liveness properties as well. This can be achieved via HML with recursion, which is again treated in modal decomposition approaches [1].

We would also like to investigate the benefits of casting causality within a process algebraic setting. Observe that, for instance, causal projections can be naturally expressed as CCS process terms derived

from CCS terms for components or their underlying LTS's. Hence, we would like to study whether a process algebraic handling of causality provide more insight on its properties and whether causality as described in this paper can be axiomatized.

Last, but not least, we would like to investigate to what extent our definition of causality is related to the actual causality in [17, 3]. As already discussed in the current paper, the two notions share similar characteristics, including causal non-occurrence of events and the ordering condition (that is implicit in our approach). Once such a relationship is identified, one could exploit the compositionality results to improve fault localisation in automated tools for causality checking [17, 3].

Acknowledgements We thank the anonymous reviewers of CREST 2016 for their constructive comments and references to the literature. The work of Georgiana Caltais was partially supported by an Independent Research Start-up Grant founded by Zukunftskolleg at Konstanz University. The work of Mohammad Reza Mousavi has been partially supported by the Swedish Research Council (Vetenskapsrådet) award number: 621-2014-5057 (Effective Model-Based Testing of Concurrent Systems) and the Swedish Knowledge Foundation (Stiftelsen för Kunskaps- och Kompetensutveckling) in the context of the AUTO-CAAS HöG project (number: 20140312).

References

[1] Luca Aceto, Arnar Birgisson, Anna Ingólfsdóttir & Mohammad Reza Mousavi (2011): *Decompositional Reasoning about the History of Parallel Processes*. In Farhad Arbab & Marjan Sirjani, editors: *Fundamentals of Software Engineering - 4th IPM International Conference, FSEN 2011, Tehran, Iran, April 20-22, 2011, Revised Selected Papers*, *Lecture Notes in Computer Science* 7141, Springer, pp. 32–47, doi:10.1007/978-3-642-29320-7_3.

[2] H. R. Andersen (1995): *Partial Model Checking (Extended Abstract)*. In: *LICS*, pp. 398–407, doi:10.1109/LICS.1995.523274.

[3] Adrian Beer, Stephan Heidinger, Uwe Kühne, Florian Leitner-Fischer & Stefan Leue (2015): *Symbolic Causality Checking Using Bounded Model Checking*. In Bernd Fischer & Jaco Geldenhuys, editors: *Model Checking Software - 22nd International Symposium, SPIN 2015, Stellenbosch, South Africa, August 24-26, 2015, Proceedings*, *Lecture Notes in Computer Science* 9232, Springer, pp. 203–221, doi:10.1007/978-3-319-23404-5_14.

[4] Mitra Tabaei Befrouei, Chao Wang & Georg Weissenbacher (2014): *Abstraction and Mining of Traces to Explain Concurrency Bugs*. In Borzoo Bonakdarpour & Scott A. Smolka, editors: *Runtime Verification - 5th International Conference, RV 2014, Toronto, ON, Canada, September 22-25, 2014. Proceedings*, *Lecture Notes in Computer Science* 8734, Springer, pp. 162–177, doi:10.1007/978-3-319-11164-3_14.

[5] G. Caltais, S. Leue & M.R. Mousavi (2016): *(De-)Composing Causality in Labeled Transition Systems*. Technical Report soft-16-02. Available at `http://se.uni-konstanz.de/uploads/tx_sibibtex/crest_2016.pdf`.

[6] D. Giannakopoulou, C. S. Pasareanu & H. Barringer (2005): *Component Verification with Automatically Generated Assumptions*. *Autom. Softw. Eng.* 12(3), pp. 297–320, doi:10.1007/s10515-005-2641-y.

[7] Gregor Gößler & Lacramioara Astefanoaei (2014): *Blaming in component-based real-time systems*. In: *2014 International Conference on Embedded Software, EMSOFT 2014*, ACM Press, pp. 7:1–7:10, doi:10.1145/2656045.2656048.

[8] Gregor Gößler & Daniel Le Métayer (2015): *A general framework for blaming in component-based systems*. *Sci. Comput. Program.* 113, pp. 223–235, doi:10.1016/j.scico.2015.06.010.

[9] Gregor Gößler, Daniel Le Métayer & Jean-Baptiste Raclet (2010): *Causality Analysis in Contract Violation*. In: *Runtime Verification - First International Conference, RV 2010*, *Lecture Notes in Computer Science* 6418, Springer, pp. 270–284, doi:10.1007/978-3-642-16612-9_21.

[10] Gregor Gößler & Jean-Bernard Stefani (2016): *Fault Ascription in Concurrent Systems*. In: *Trustworthy Global Computing - 10th International Symposium, TGC*, *Lecture Notes in Computer Science* 9533, Springer, pp. 79–94, doi:10.1007/978-3-319-28766-9.

[11] A. Groce, S. Chaki, D. Kroening & O. Strichman (2006): *Error explanation with distance metrics*. *International Journal on Software Tools for Technology Transfer (STTT)* 8(3), doi:10.1007/s10009-005-0202-0.

[12] A. Groce & W. Visser (2003): *What Went Wrong: Explaining Counterexamples*. In: *Workshop on Software Model Checking (SPIN)*, Lecture Notes in Computer Science 2648, Springer, pp. 121–135, doi:10.1007/3-540-44829-2_8.

[13] J.Y. Halpern & J. Pearl (2005): *Causes and explanations: A structural-model approach. Part I: Causes*. The British Journal for the Philosophy of Science, doi:10.1093/bjps/axi148.

[14] Matthew Hennessy & Robin Milner (1980): *On Observing Nondeterminism and Concurrency*. In J. W. de Bakker & Jan van Leeuwen, editors: Automata, Languages and Programming, 7th Colloquium, Noordweijkerhout, The Netherland, July 14-18, 1980, Proceedings, Lecture Notes in Computer Science 85, Springer, pp. 299–309, doi:10.1007/3-540-10003-2_79.

[15] Peter Ladkin & Karsten Loer (1998): *Analysing Aviation Accidents Using WB-Analysis – an Application of Multimodal Reasoning*. In: AAAI Spring Symposium, AAAI. Available at `https://www.aaai.org/Papers/Symposia/Spring/1998/SS-98-04/SS98-04-031.pdf`.

[16] Kim Guldstrand Larsen & Liu Xinxin (1991): *Compositionality through an Operational Semantics of Contexts*. J. Log. Comput. 1(6), pp. 761–795, doi:10.1093/logcom/1.6.761.

[17] Florian Leitner-Fischer & Stefan Leue (2013): *Causality Checking for Complex System Models*. In Roberto Giacobazzi, Josh Berdine & Isabella Mastroeni, editors: Verification, Model Checking, and Abstract Interpretation, 14th International Conference, VMCAI 2013, Rome, Italy, January 20-22, 2013. Proceedings, Lecture Notes in Computer Science 7737, Springer, pp. 248–267, doi:10.1007/978-3-642-35873-9_16.

[18] Florian Leitner-Fischer & Stefan Leue (2013): *Probabilistic Fault Tree Synthesis using Causality Computation*. International Journal of Critical Computer-Based Systems 4, pp. pp. 119–143, doi:10.1504/IJCCBS.2013.056492.

[19] Florian Leitner-Fischer & Stefan Leue (2014): *SpinCause: a tool for causality checking*. In Neha Rungta & Oksana Tkachuk, editors: 2014 International Symposium on Model Checking of Software, SPIN 2014, Proceedings, San Jose, CA, USA, July 21-23, 2014, ACM, pp. 117–120, doi:10.1145/2632362.2632371.

[20] D. Lewis (1973): *Counterfactuals*. Blackwell Publishers.

[21] Robin Milner (1980): *A Calculus of Communicating Systems*. Lecture Notes in Computer Science 92, Springer, doi:10.1007/3-540-10235-3.

[22] N. Noroozi, M.R. Mousavi & T.A.C. Willemse (2013): *Decomposability in Input Output Conformance Testing*. In: Proceedings of the 8th Workshop on Model-Based Testing (MBT 2013), Electronic Proceedings in Theoretical Computer Science 111, pp. 51–66, doi:10.4204/EPTCS.111.5.

[23] M. Renieris & S.P. Reiss (2003): *Fault localization with nearest neighbor queries*. In: 18th International Conference on Automated Software Engineering, Montreal, Canada, doi:10.1109/ASE.2003.1240292.

[24] T. Villa, N. Yevtushenko, R.K. Brayton, A. Mishchenko, A. Petrenko & A. Sangiovanni-Vincentelli (2012): *The Unknown Component Problem, Theory and Applications*. Springer, doi:10.1007/978-0-387-68759-9.

[25] G. Xie & Z. Dang (2006): *Testing systems of concurrent black-boxes—an automata-theoretic and decompositional approach*. In: FATES, LNCS 3997, Springer, pp. 170–186, doi:10.1007/11759744_12.

[26] Andreas Zeller (2009): *Why Programs Fail: A Guide to Systematic Debugging*. Elsevier.